

# End-to-End Evaluation of Amazon's S3 and EC2

Simson L. Garfinkel



# Amazon's Web Services

- Originally: E-Commerce, Alexa,
- Summer 2006: Commodity Computing
  - S3 — Simple Storage Service
  - EC2 — Elastic Compute Cloud
  - SQS — Simple Queue Service
- BusinessWeek: "Jeff Bezo's Risky Bet"



# Why use these services?

- Backed by one of the world's largest clusters.
- Highly redundant, highly scalable.
- Cheaper than building it yourself.
- Isolates from the "slashdot effect."



# S3: "Unlimited" Storage

- 10 ¢/GB/month for storage
- 15 ¢/GB to upload or download
- SOAP or "REST"
  - <http://s3.amazonaws.com/bucket/object>
  - <http://bucket.s3.amazonaws.com/object>
- Object size 1-5GB (2GB); 4KB Metadata



# EC2: "Unlimited" Compute

- 10 ¢/CPU-HOUR
- Xen-based Virtual machine with:
  - 1.7 Ghz x86 processor
  - 1.75 GB RAM
  - 160 GB hard drive
  - 250 MB/s of network bandwidth



# SQS: Easy Queues

- 10 ¢/10,000 messages
- 0-256K message size with metadata
- Unlimited Queue Size
- Reliable Delivery; Random Order
- PUT/READ-LOCK/DELETE functionality
- SOAP & REST API



# Who is the customer for these services?

- Amazon won't say. "They're for everybody!"
  - Large customers that need reliable off-site backup.
  - E-businesses that need scalable infrastructure.



# My Analysis: Summary

- The services work... most of the time.
- The services scale well.
- Key problem points:
  - Availability
  - Security model
  - Lack of an SLA



# Experimental Setup



25,066 probes between Nov. 3 and Jan 3rd.



# Experimental Setup 2



94,615 probes between Nov. 3 and Jan 3rd.



# Average Performance (EC2)

Size	GET	PUT
1B	40 TPS	19 TPS
1M	1 MBS/sec	4 MB/sec
100M	9 MB/sec	6 MB/sec



# Peak (top 1%) Performance (EC2)

Size	GET	PUT
1B	76 TPS	42 TPS
1M	1 MBS/sec	8 MB/sec
100M	19 MB/sec	14 MB/sec



# Best performance: when routers failed

- GET performance: 21 MBytes/sec
- PUT performance: 17 MBytes/sec
- Success rate: <5%

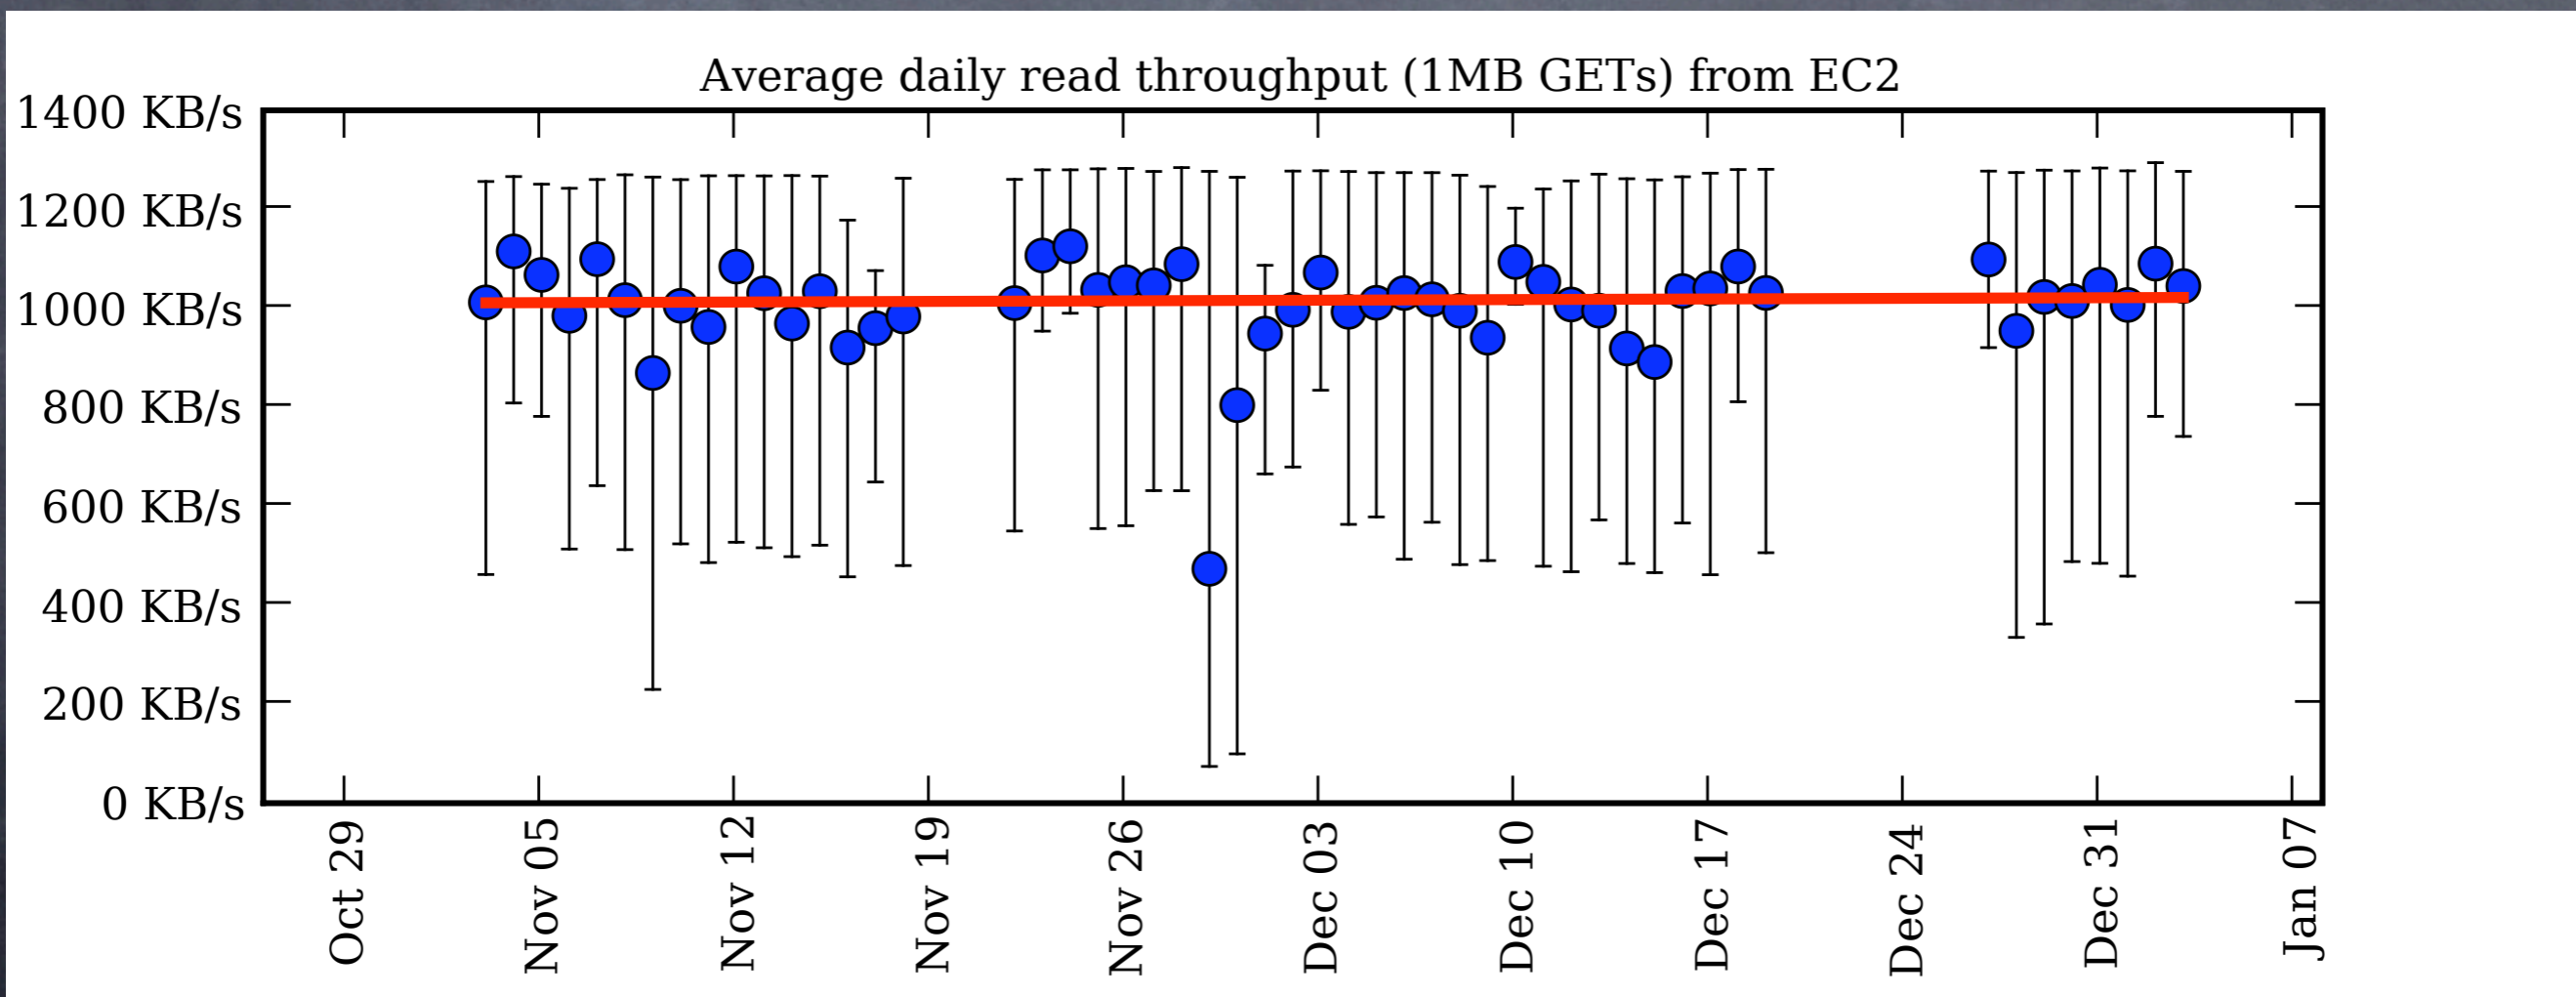


# Performance from outside Amazon (KB/sec)

Host	avg GET	avg PUT
EC2	1024	3760
Harvard 1	741	1360
Harvard 2	638	847
MIT 1	726	1730
MIT 2	796	2077
ISP PIT	599	1428
ISP LA	930	933
FiOS	548	199

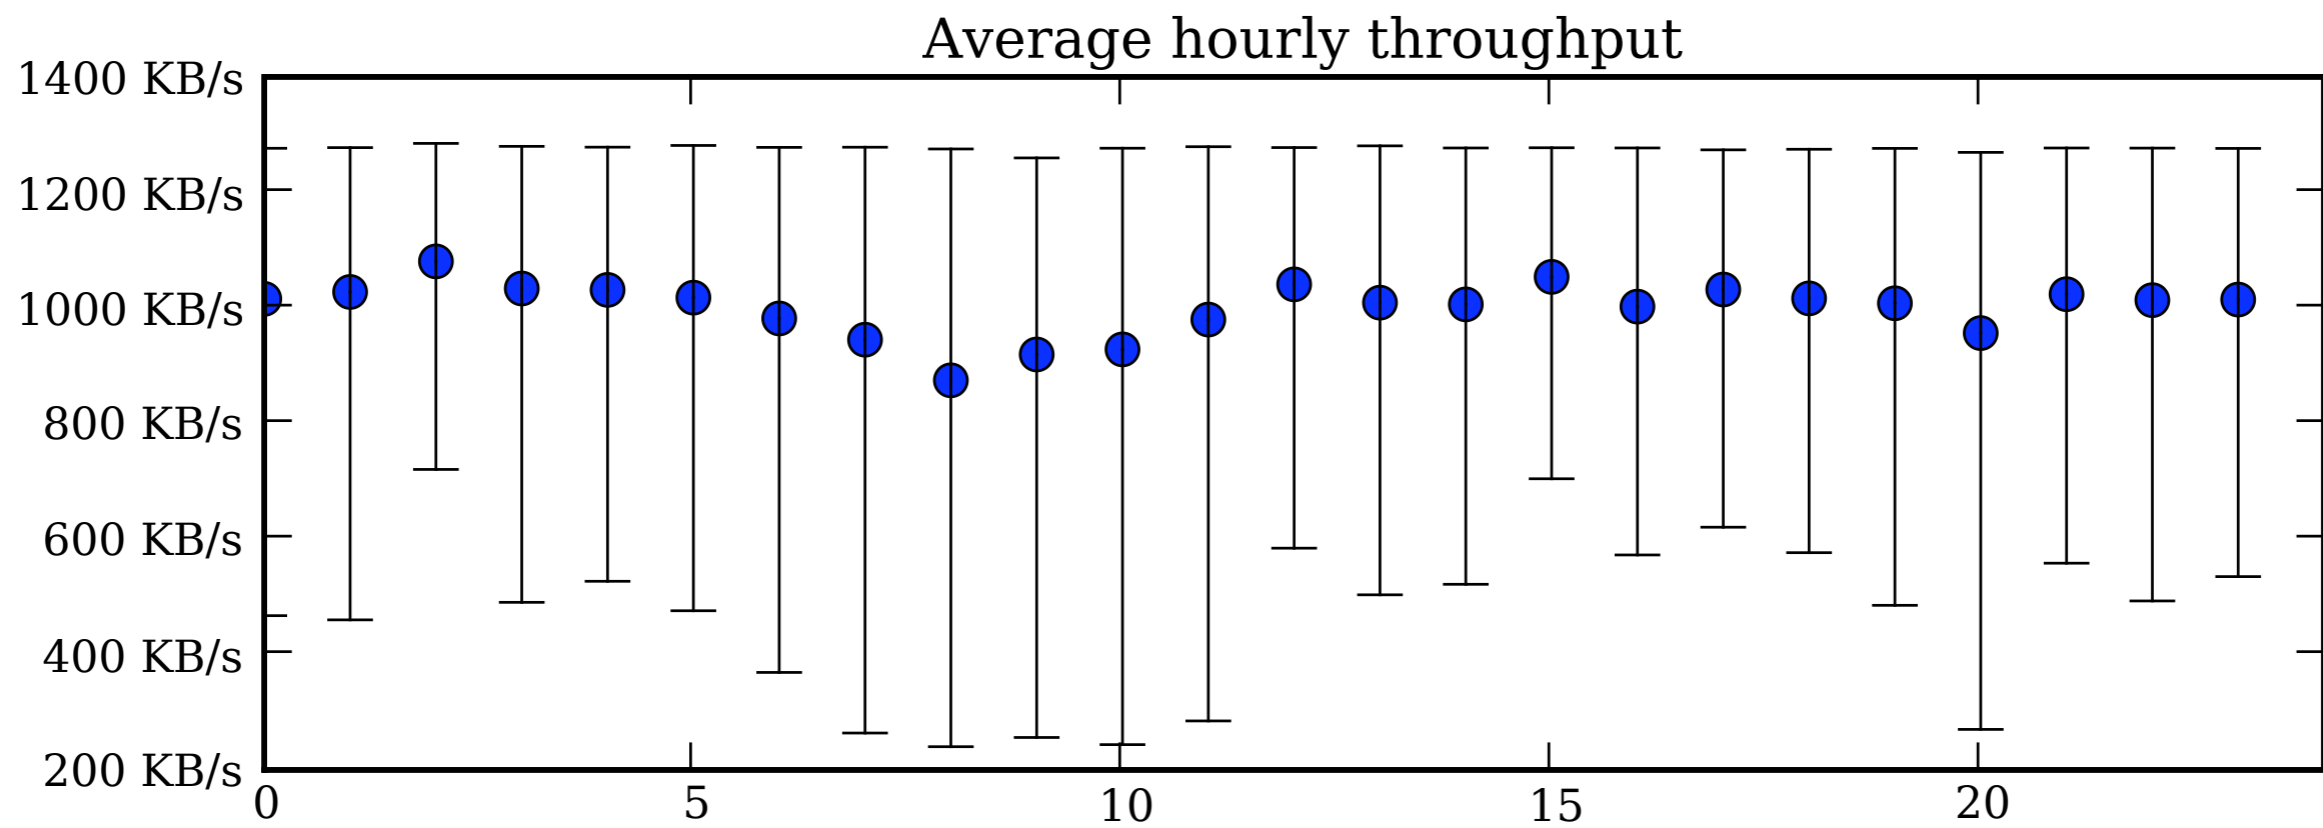


# Average daily read throughput (1MB GETs)



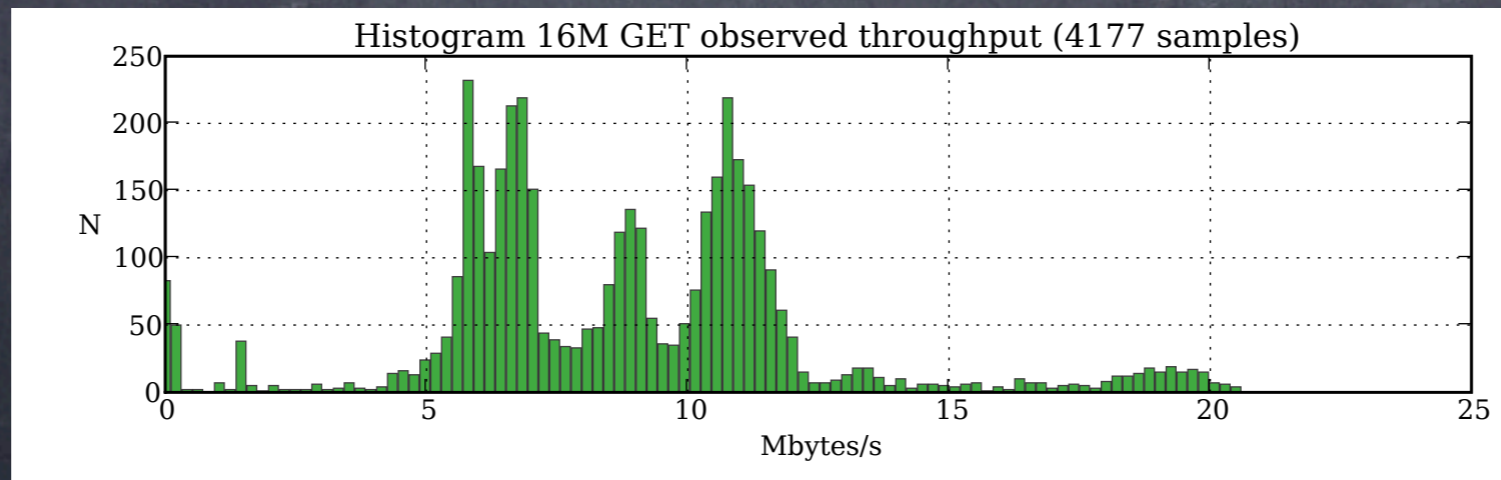
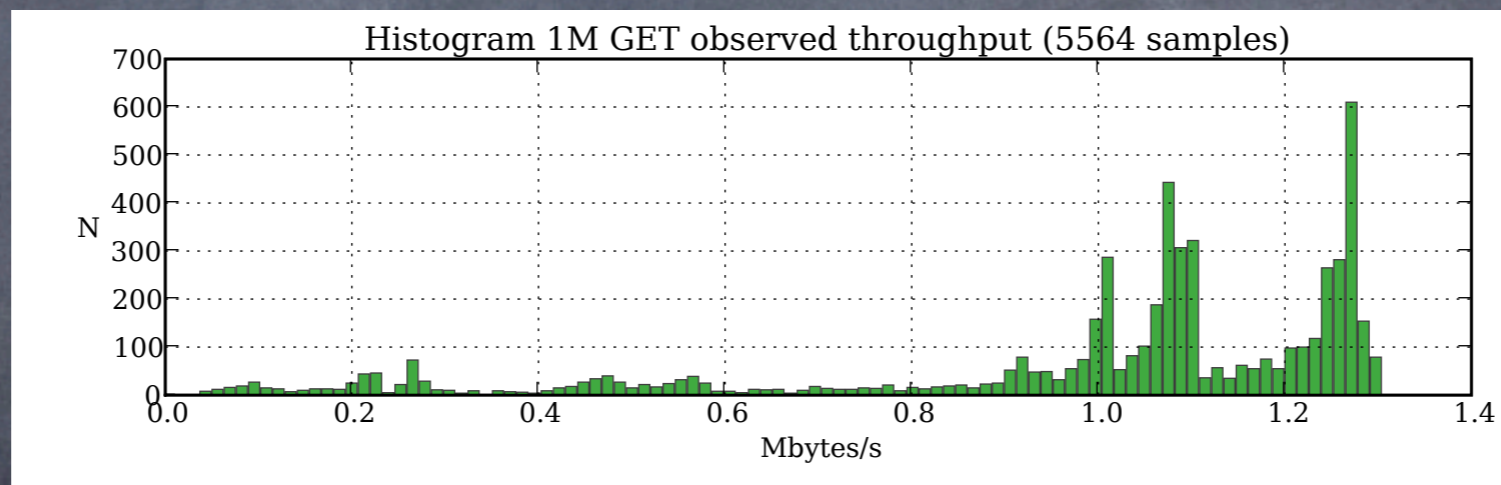
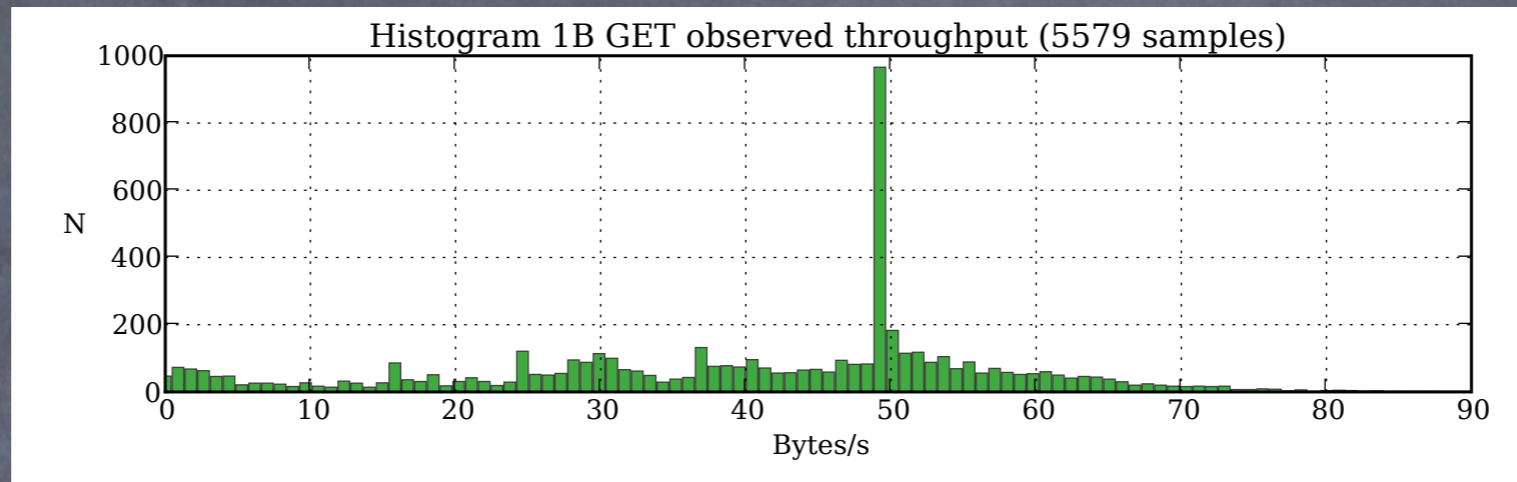


# Average hourly throughput




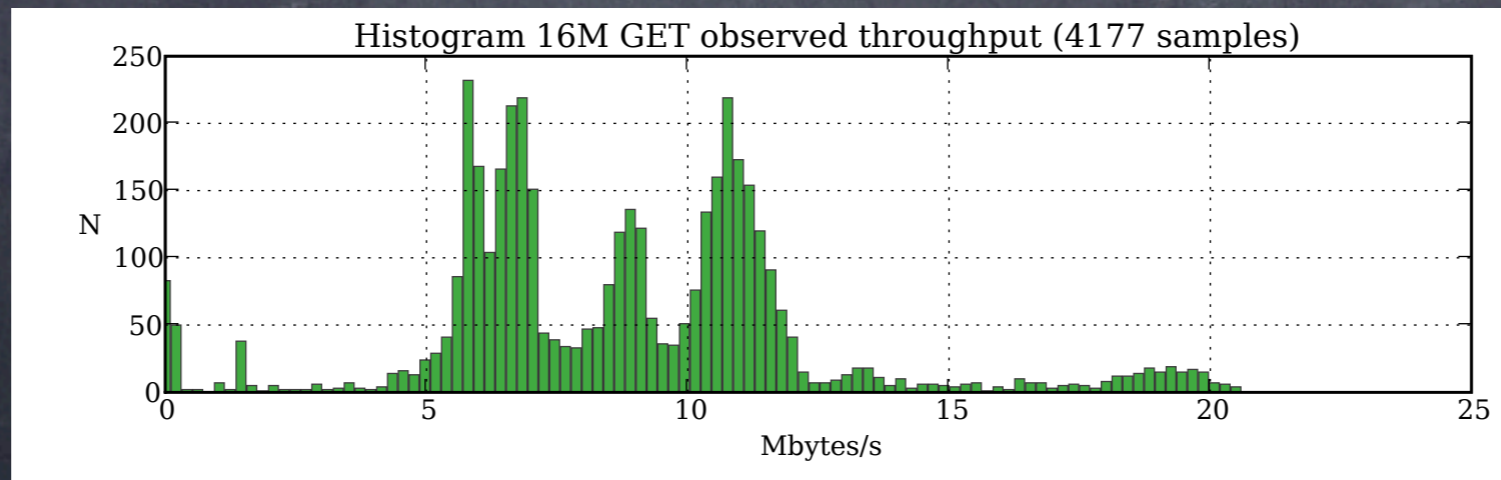
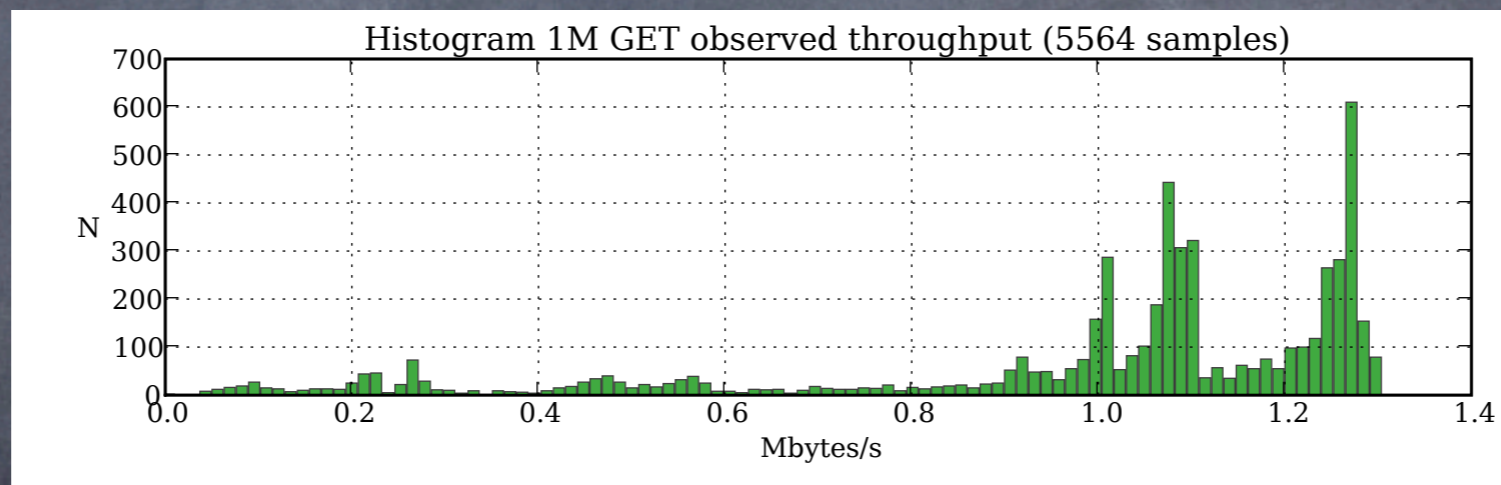
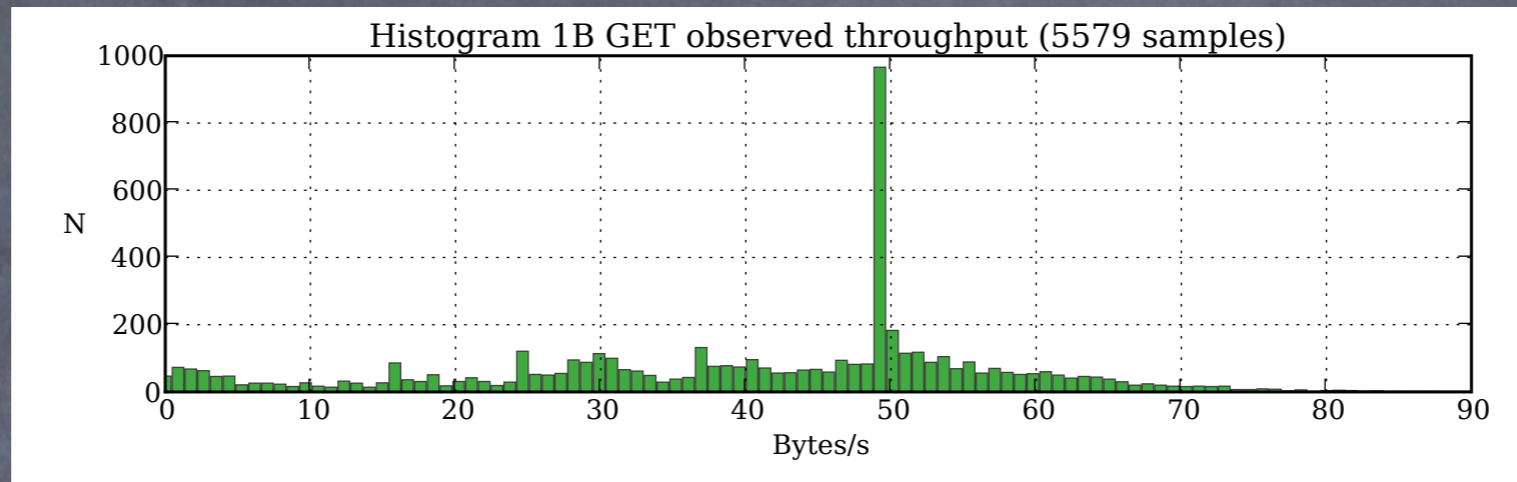


# GET Distribution





# GET Distribution



3



# Security Issues 1: Authentication



# Security Issues 1: Authentication

- Per-Object & Per-Bucket access control lists



# Security Issues 1: Authentication

- Per-Object & Per-Bucket access control lists
- Authentication with HMAC & Secret Key



# Security Issues 1: Authentication

- Per-Object & Per-Bucket access control lists
- Authentication with HMAC & Secret Key
- Secret Key downloaded from AWS website



# Security Issues 1: Authentication

- Per-Object & Per-Bucket access control lists
- Authentication with HMAC & Secret Key
- Secret Key downloaded from AWS website
- Password reset with email challenge/response



# Security Issues 1: Authentication

- Per-Object & Per-Bucket access control lists
- Authentication with HMAC & Secret Key
- Secret Key downloaded from AWS website
- Password reset with email challenge/response
- Control of email = control of all S3 data



# Security Issues 2: availability

- S3 has experienced repeated service interruptions.
- Amazon's long-term commitment to the service is unknown.



# Security Issues 3:

## Amazon Web Services License Agreement

- Allows termination at any time for any reason.
- Prohibits content that is "obscene, libelous, defamatory or otherwise malicious or harmful to any person or entity"
- Prohibits service "in any way that is otherwise illegal or promotes illegal activities, including without limitation in any manner that might be discriminatory based on race, sex, religion, nationality, disability , sexual orientation or age."